Data Intensive Systems

7th PI Meeting

New Orleans, LA

May 17-18, 2001

Robert B Graybill

(703) 696-2220

Fax: (703) 696-4534

email: rgraybill@darpa.mil

DIS Motivation

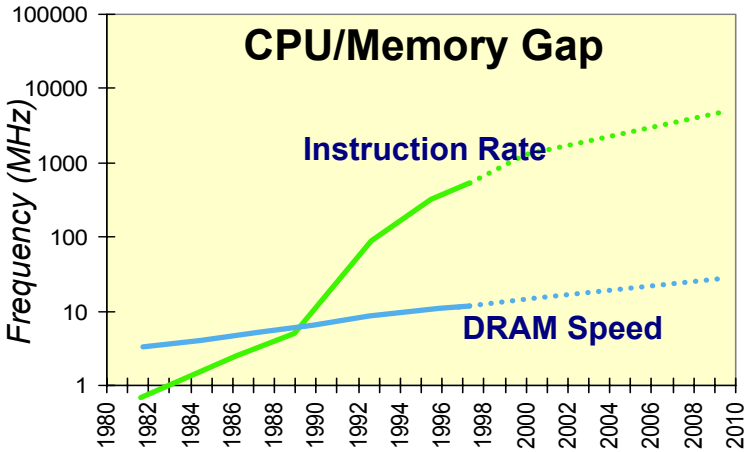

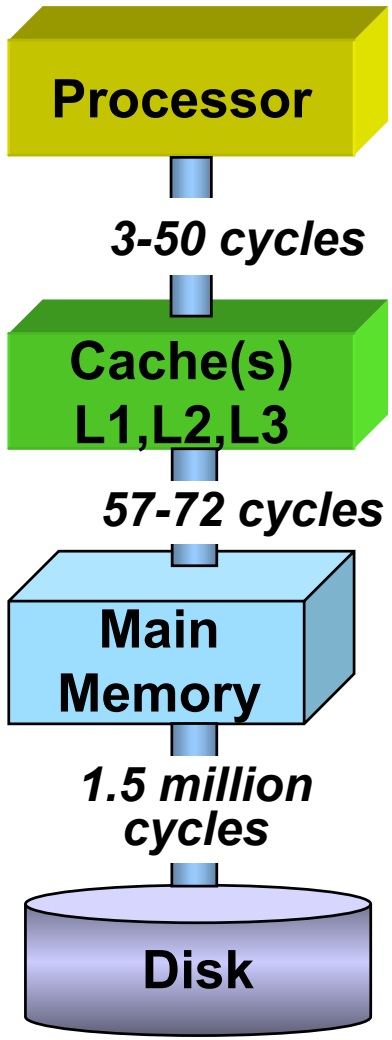

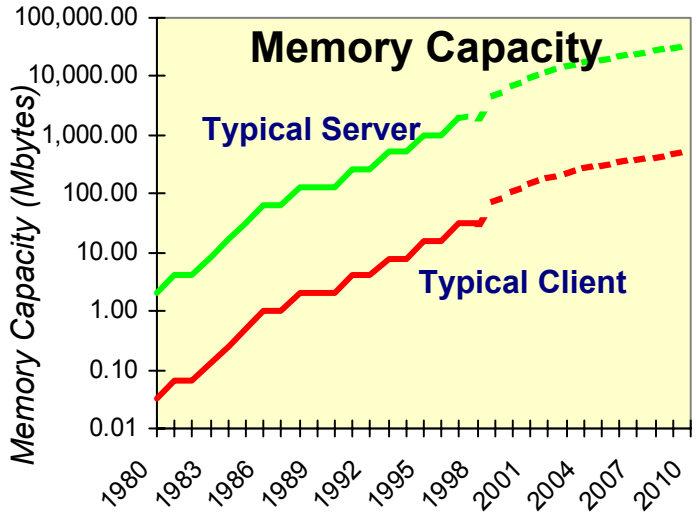

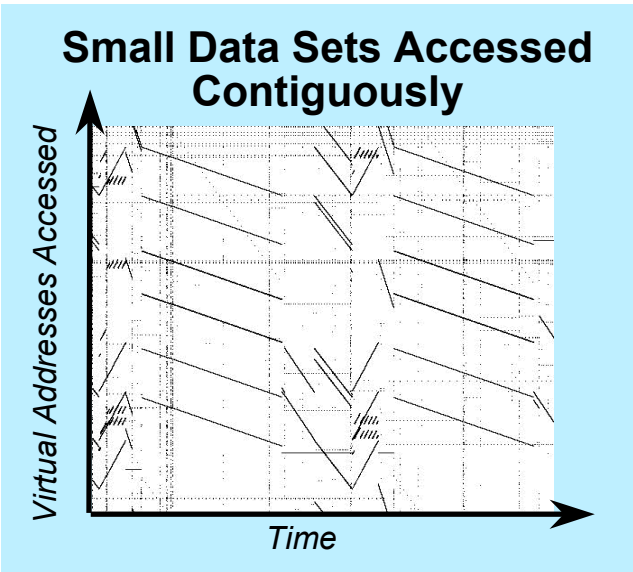

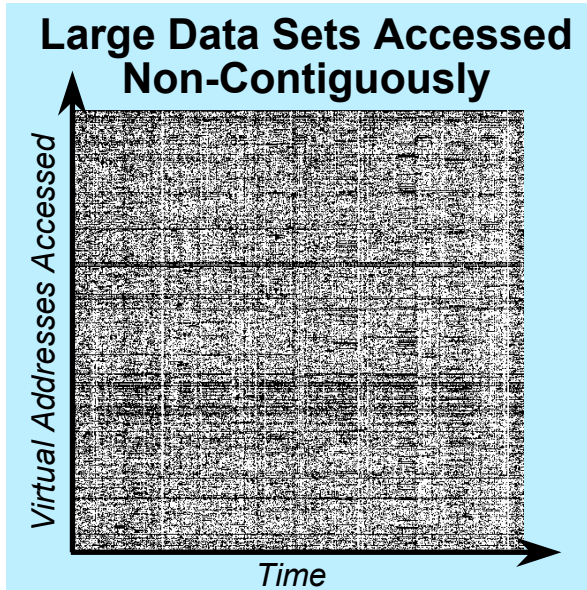

Memory access lags memory growth and processor performance

DIS Program Application Tenets

Data Intensive Applications Are:

- Large high-rate data streams (data-rate limited)
- Distributed data access (non-contiguous data)
- Dynamic data accesses (lack of access predictability)

Data Intensive Applications Are NOT:

- Applications with small working sets of contiguous data

Approach

In Situ Processing

Processor within Memory

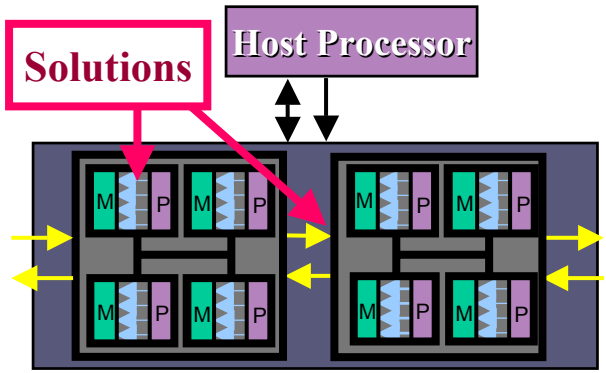

Adaptive Cache Management

Adaptive Memory

Leverage application knowledge and run-time information to extract locality from apparently pseudorandom access patterns.

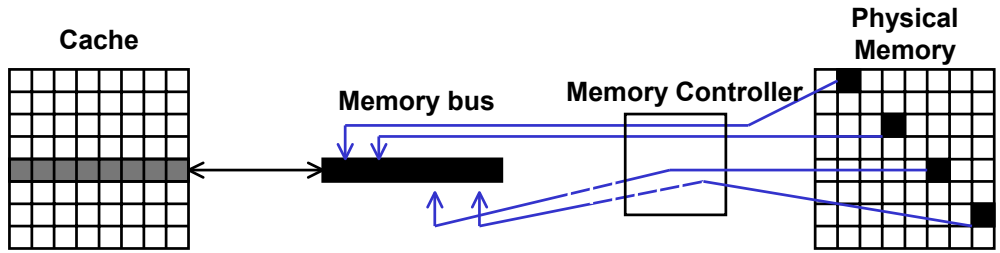

Computation Stream within Memory

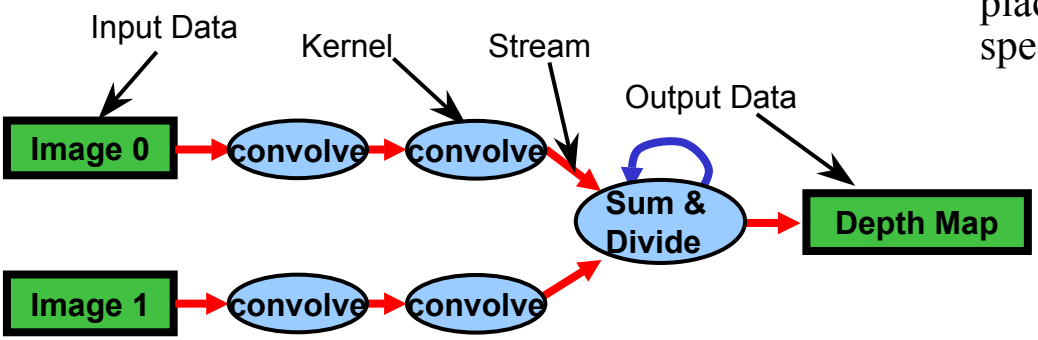

Algorithm, Compile & Data Placement

Applications manage memory hierarchy so data placement and flow is tailored to application specific needs.

Perfect Latin Squares: conflict-free access

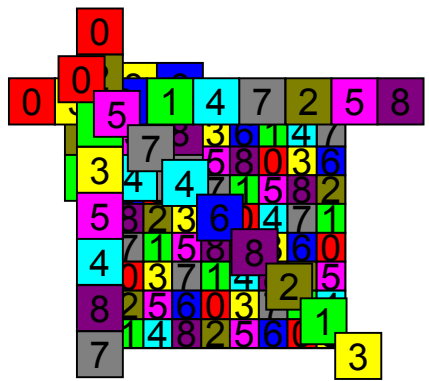

DIS Projects Matrix

AO#	Project Title	Organization	In Situ Processing		Adaptive Cache Management		System Architecture
			PIM	CSIM	AM	AC & DP	
D347	FLASH Graphics System & Architecture	Stanford University					◆
E254	ImagEn: Image Manipulation Engine	Stanford University		◆	◆		
E336	Intelligent DRAM (IRAM) and ISTORE	University of California, Berkeley	◆			◆	◆
E393	Adaptive Structure Aware Memory Systems	University of Utah			◆	◆	
G182	Adaptive Memory Reconfiguration Management (AMRM)	University of California, Irvine			◆	◆	
G183	Active Database Technology	Massachusetts Institute of Technology			◆	◆	◆
G185	Smart Memories: A Universal Computing Element	Stanford University	◆	◆	◆	◆	◆
G215 *J099	DIVA	USC Information Sciences Institute	◆			◆	
H306	SLIIC	USC Information Sciences Inst. - East	◆				◆
H307	PIM Applications & DIS Benchmarks	Atlantic Aerospace Electronics Corp.					◆
H309	Algorithms for Data Intensive Applications on Intelligent & Smart Memories (ADVISOR)	University of Southern California			◆	◆	
H316	Malleable Caches for Data-Intensive Computing	Massachusetts Institute of Technology			◆	◆	
H318	Algorithmic Strategies for Compiler Controlled Smart Caches	New York University			◆		

PIM – Processor within Memory
CSIM – Computation Stream within Memory
AM – Adaptive Memory
AC & DP – Algorithm, Compile, & Data Placement

DIS Benchmarks and Stressmarks Correlation Matrix

Benchmarks are focused code that retain the context of the enveloping application.

Stressmarks are a collective suite of smaller, specific procedures that illustrate DIS attributes.

	SAR Ray-Tracing	Data Mgmt	Multidimensional DFT	Image Understanding
Pointer Stressmark Small blocks at unpredictable locations. Can be threaded.	□	■		
Update Stressmark Small blocks at unpredictable locations.	□	■		
Matrix Stressmark Dependent on matrix representation. Likely to be irregular or mixed, with mixed densities.	□		□	□
Neighborhood Stressmark Regular access to pairs of words at arbitrary distances, and irregular access to histogram.		□	□	■
Field Stressmark Regular, dense, little re-use.		■		□
Corner-Turn Stressmark Data movement w/o processing.	□		■	□
Transitive Closure Stressmark Dependent on data organization. Likely to be dense and irregular or mixed.	□			■

□ = " Correlated": the benchmark contains an operation similar or identical to the operation that the stressmark emphasizes.

■ = "Strongly correlated": the operation is present, and further, it is a major contributor to the performance of the benchmark.

Program Schedule

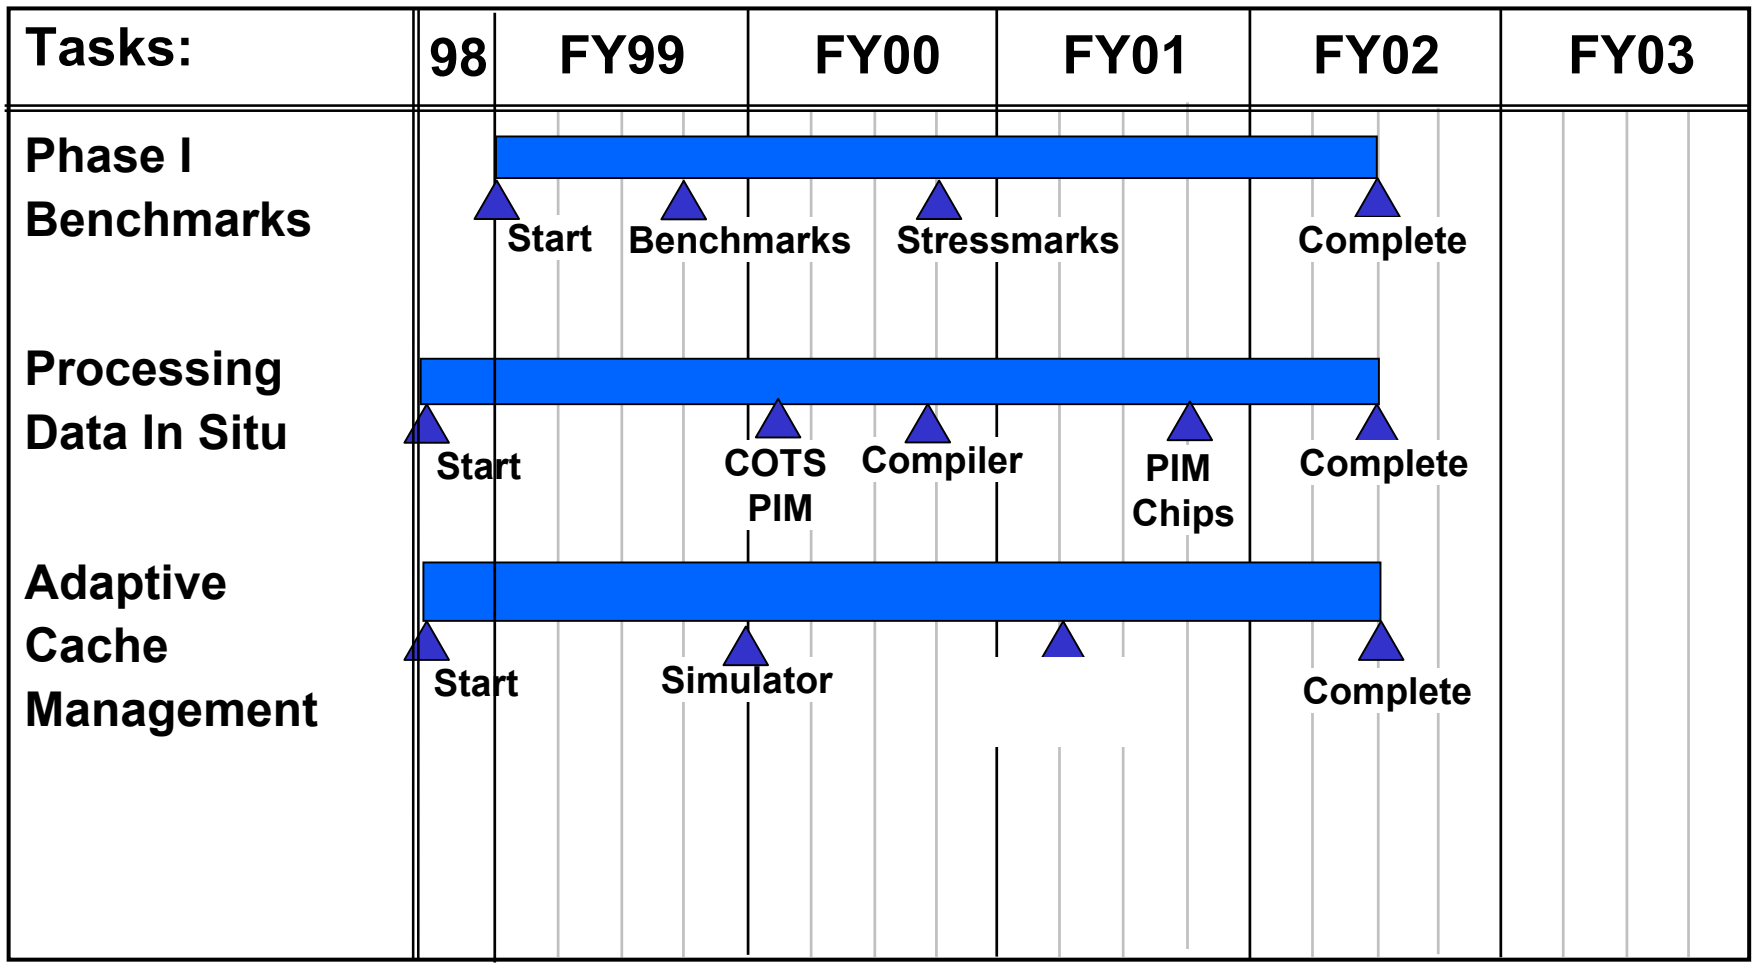

DIS Technology Transfer

FLASH Graphics & Imagine Chip

Origin 2000

Imagine Chip

Texas Instruments

Vectorizing Compiler

CRAY

ISTORE

Anigma Inc.

VIRAM

IRAM Scalar Core

MIPS TECHNOLOGIES

DIVA

PIMs for Blue Gene

IBM

USC INFORMATION SCIENCES INSTITUTE

EAST

SLIIC PIM for distributed beam steering array control payload & coherent sidelobe canceller

LOCKHEED MARTIN

Navy Aegis

Georgia Institute of Technology

Compiler and Tools for Explicitly Parallel Instruction Computing (EPIC)

hp invent

intel

THE UNIVERSITY OF UTAH

Impulse

Model

hp invent

Sun microsystems

AMRM Board and tools

UCI University of California, Irvine

C++ Based Specification Language

tensilica

SYNOPSYS

Memory Models

FPGAs

denali Memory solutions defined.

ALTERA